HIT Standards: Privacy & Security Workgroup
Draft Transcript
April 23, 2010

Presentation

Judy Sparrow – Office of the National Coordinator – Executive Director
Good afternoon, everybody, and welcome to the Privacy & Security Workgroup. This is a federal advisory committee. There will be opportunity at the end of the call for the public to make comments. Let me do a quick roll call. Dixie Baker?

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Anne Castro?

Anne Castro – BlueCross BlueShield South Carolina – Chief Design Architect
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Steve Findlay? Dave McCallie couldn't make it. Gina Perez? Wes Rishel? Sharon Terry? Walter Suarez?

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Yes, I'm here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Aneesh Chopra? Jodi Daniel? Joy Pritts? Deborah Lasky? John Moehrke?

John Moehrke – Interoperability & Security, GE – Principal Engineer
Present.

Judy Sparrow – Office of the National Coordinator – Executive Director
Sue McAndrew? Sarah Wattenberg?

Sarah Wattenberg – ONCHIT – Public Health Advisor
Yes.

Judy Sparrow – Office of the National Coordinator – Executive Director
We did invite members of the privacy and security policy workgroup. I know we've got Deven McGraw, Judy Faulkner, John Blair, Paul Uhrig, and John Houston on the line. Anybody else from …?

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
Kathleen.

Judy Sparrow – Office of the National Coordinator – Executive Director
Kathleen. That's it. Dixie, I can turn it over to you.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences

Okay. First, I want to apologize for dialing in late. I got to doing some real work, and the time just slipped up. This is the second of a series of educational sessions that we've scheduled for the privacy and security standards workgroup. For the purpose of our getting a good, basic understanding of the standardization efforts that have been done, and that are in progress that are relevant to the management of consumer consent. And in each of these sessions, I want to start out by giving you the basic context for this discussion.

The first, there are different kinds of consents that we deal with. The two basic consents or permissions that are relevant to HIPAA and the privacy rule of HIPAA are the basic privacy consents, which are either explicit or implicit, which can be the consumer's written or verbal permission to collect, use and disclose individually identifiable health information and privacy authorization, which is always written, and it's a signed, written document that contains explicit elements that are set forth in the HIPAA privacy rule and gives the covered entity permission to use or disclose the individually identifiable health information for specific purposes. Examples of these are like for release for research, release for marketing, those kinds of things.

There's also, under the common rule and FDA rules and other HHS rules, the concept of informed consent. These are consents that still have to be captured in the electronic health record, and they're written permissions to perform a particular activity or for the patient to participate in a clinical trial, etc. When we talk about consent, we're talking about them broadly, and we're addressing different standards that relate to those.

Now historically, we've managed consent as just paper, paper records, written. The written authorizations for HIPAA, as well as informed consent have to be signed by the individual patient, so there is a requirement for a signature there. In the future, however, these signatures are likely to be digital. So what will happen in the future is that the patient or consumer would digitally sign the consent or authorization. Then somehow those permissions have to be captured as part of the health record. That isn't to imply that physically they have to reside in the EHR system, but the health record needs to be able to link to these permissions when necessary, so we really aren't discussing where they're persisted or where they're kept over time.

The permissions have to be interpretable by humans, as well as computers. Today, most permissions are interpretable by humans, but not generally by computers. So they also have to be cross-validated. You can easily conceive of a patient giving permissions that could be inconsistent between the two, so this computer has to figure out what are the valid permissions from all of the permissions that are available to it.

Then somehow those rules have to be tied to the information that is exchanged. As it's sent over the NHIN, for example, it needs to be tied to that information, and somehow the user's permission, the permissions that were granted, need to be persisted with the data somehow over time. Now we do lack specific policy in this area right now, but this is conceptually what a consumer is likely to expect in the future.

The standards that are needed here are in the area of digital signatures. Privacy policies that include state policy, federal policy, local policy, all of those together, and again, the system has to be able to identify any inconsistencies among the two or duplications and decide what policies to actually apply. A data model and a schema, permission syntax, how you represent those permissions in a system such that the computer can understand it within a particular context, and vocabulary, just like we talk a lot at

the standards committee about vocabularies for capturing clinical symptoms and clinical orders, laboratory results. There need to be vocabulary for permissions as well.

Then there need to be standards for how we do cross validation and how we resolve inconsistencies, how we maintain permissions and retrieve permissions within an organization or between organizations, and then how those permissions are translated into access control rules. If we exchange a permission, let's say we have permission X in my hospital, and I send it over to Judy Sparrow, and she goes, "Oh, permission X. John Moehrke has permission X." If we don't both interpret that permission X in the same way, then the rules will not be enforced in the same way across organizations or even across systems within a single organization. So that translation between the consent that's given or the permission that's given and how it's interpreted and translated into access control rules is really important. Then the whole enforcement and auditing of what happens and how these permissions are enforced needs to be captured.

Then, finally, how we exchange permissions between organizations within an HIE or an NHIN, and how permissions are updated and revoked. If I decide that I no longer want my health information to be used for research, let's say, I no longer want it to be used, how does that get propagated across an HIE or an NHIN, or is once it's out there, is everything lost? All of these things are areas that we need to really look for standards to cover.

Today, we are going to be talking about the basic patient privacy consent profile, which was developed by the integrating the healthcare enterprise, the IHE, initiative. And John Moehrke has agreed to talk to us about this. John has very, very deep experience with BPPC and with IHE in general, so I'm very grateful to him for agreeing to present this today. As you'll see, the BPPC covers a signature capture, whether it be a wet signature or a digital signature. It also has a data model and a schema inside and a certain level of permission syntax and vocabulary, not exhaustible, but some vocabulary. Then it also addresses the exchange and sharing of permissions across systems and across and between enterprises.

With that, I'll turn this over to John, or I hope someone will turn this over to John.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Thanks, Dixie. Yes. I'm John Moehrke. For those on the HIT Standards Privacy & Security Committee, you know me well as a member. I welcome the policy members as well. I am a GE Healthcare employee, but in this context here, I'll be speaking more as a member of the IHE committee that works on the basic patient privacy consent profile, so that's why the color changes, and the presentation is more IHE centric, as this is a reusable presentation from IHE.

IHE started to look at how we can support consent in interoperability profiles back in 2006. We had a meeting in New Brunswick, spent a couple of days trying to find standards that we could profile as IHE for the use of this. Essentially, as Dixie showed quite nicely, there is so much that you need to have in place in order to truly capture the needs of a consent or even an authorization and such.

In the absence of these standards, but still having the need, we decided we have to do something basic, so the very beginning of this profile name, basic, really is a key name in what we provided. And we knew, right out of the gate, back four years ago that it was not going to be the ultimate solution. It was not necessarily even going to be a minimalistic support, but we had to provide some form of basic support for the capturing and communications of consent.

As Dixie also explained, consent is made up of many different things. Indeed, there is the policy aspect, which is really the part that is in many cases very hard to do, trying to write the policies that truly express

how the data is expected to be moved, how that will work within the workflows of the organization. It needs to be very risk based so that you look at what are the risks. Indeed, any time you are moving patient data, you're at least introducing a risk to the patient of their data being exposed, but you're also introducing risks to the business. You're introducing risks to the information being disclosed and such.

These policies are made up of what are the rights and the responsibilities, what are the acceptable uses, and where can this data go, and what needs to be done when the data goes there. It also really needs to include, well, what is the result if the policies are not followed? And oftentimes that's not necessarily as obvious in the policy, but when you're dealing with a basic patient privacy consent, we knew that it was very important to expressly indicate that because this is basic, there probably will be plenty of edges that are not as well controlled and may have to rely more on policy that essentially might allow the technology to do something that is not allowed by policy. If that's the case, you need to do this post access enforcement, if you will.

The other aspect of this is what is the process by which the consent is dealt with? How are the policies enforced? Who are the people that are involved? What are the organizations that are authorized, essentially, the good old, who, what, when, where, and how? And then you finally get to what is the technology that's doing this enforcing? What is the technology that is capturing the consent? What are the algorithms that are used? This gets down to the usual kind of stuff that we like to gravitate towards as technical folks.

But essentially, we have to recognize that there's really far more than just the technology involved and truly, the IHE, as an organization, really only focuses on specifying the interoperability aspect of the technology. It doesn't specify the thing that would be used to capture the ink on paper or that kind of technology. It only captures, well, what would be the result of how you would communicate so that there is an ink on paper signature.

Back in 2006, IHE had this profile called XDS, cross enterprise document sharing, and it was essentially a profile that was setting up an health information exchange, and we knew that we needed to support consent in that environment. Oftentimes, prior to this cross enterprise world, consents were always dealt with within an organization as mostly a paperwork piece where, at the registration desk, you would fill out all of your paperwork, and if you properly filled out all your paperwork, the clerk went and registered you into the workflow, that was as far as the technology cared about whether you had consented or not. But when you start to talk about cross enterprise communications of documents, you have this concern that the patient may have consented at one organization in a way that would allow, or maybe not, another organization to view the data that was in the health information exchange. Clearly this becomes a high priority, and hence why IHE spun up this program.

At the time, the documentation for XDS indicated that there was really one affinity domain for health information exchange. Therefore, there was only one policy. It was up to them to define what the policy is. It wasn't an IHE declaration. And essentially, because there's only one policy, because there is no way to capture multiple flavors of consent that a patient might give, you really ended up with these rules that said, well, gee, if the patient doesn't agree to the one policy, then the only thing you can do is not publish data into the health information exchange for that patient, which meant, when you have a VIP patient, a sports figure, your only option is to just not publish because there is nothing that you can do to say this is a highly sensitive patient, which also would be true of a patient that is sensitive from the point of view of their health status or social status, victim of violence, those kinds of things. Essentially, the results back then in the absence of any form of consent mechanism was that there were so many individuals that really were given no choice, and there were so many more individuals who really, you know, you just couldn't even use the system because it couldn't support their needs.

What IHE did was it looked at what should a consent support, and weighed that against what was available. Dixie covered a lot of these on her slides, but clearly a consent has to be able to have some human readable component to it, especially in the basic form because, in the basic form, we might actually have to leverage the consuming organization reading the text of the consent in order to enforce that particular consent. But essentially that was considered an important part that the consent be able to carry some human readable component.

It was also very important that there is a machine processable component to it. If you can't differentiate an opt in from an opt out at the machine level, you're not going to be automating anything. So we needed to be able to automate some forms of consent choice, so there needed to be a machine processable aspect to this.

We really liked the characteristics a CDA document has, and we'll cover that. We needed to be able to support multiple consent types. As Dixie pointed out earlier, there is the classic consent, but there's also a HIPAA authorization, which we included in the scope of what BPPC could or should be able to cover, so an authorization would be another form of consent. Indeed, over time, we've changed the way that we describe these things to talk more about a privacy policy acknowledgment as opposed to a consent because the word consent seems to carry some political baggage.

As Dixie pointed out today, we often capture a consent from a patient through putting the ink onto paper. Well, ink on paper is an important aspect of capturing consent. Indeed, the actual ceremony of signing has a … reaction. Indeed, there's been plenty of case law where the ink on paper was brought back out and placed in front of an individual that was claiming that they had never signed something, and truly the visual of seeing their ink on paper is sufficient for them to say, "Oh, wait a minute. Now I remember signing it." So this ritual of signing is very important, and it might continue to be important for a long time. But we clearly wanted this consent mechanism to support that.

On the other hand, IHE is an international organization, and we know of places like Germany where digital certificates are being distributed to consumers, and a consumer very well might want to or the system might want to allow for signing the consent digitally as opposed to ink on paper or both. So we also wanted to be able to support digital signatures. Then this aspect that we knew this was basic meant that we wanted to have some kind of a concept of how could this be extended as time goes by.

This is a graphic that is more contemporary. On the bottom, you see XDS. But you also see that XDR is there, XDM is there, and XDA. Those are three other ways in which documents can be communicated across enterprise boundaries. In the case of XDS, it's a shared environment where you register that you have a document available, and you do that … and 20 years later, somebody can discover that it's there and pull it. In the case of XDR, it's a point-to-point push. So if I am a provider wanting to send a referral to somebody else, I could push the documents, as a bundle of documents in a Web services document, using XDR.

The XDM environment is for media exchange. The two primary medias are CD-ROM or USB memory stick, so the case, many PHRs years ago were these memory sticks that the patient would carry around with them. Being able to lay documents on there in a way that is understandable, both for a human with just a simple browser, but also for an EHR that would need to be able to important these documents is very important. But the XDM profile also supports the delivery of a set of documents through an e-mail exchange, so there's that exchange.

Then the cross-community access is a model for how you would federate a set of XDS or other health information exchanges that are not XDS based. So we now have this infrastructure of how to move documents, and we have expanded the BPPC profile to fit within there. Essentially, you'll notice that BPPC is up in this document content profile group along with things like the scanned document or a medical summary or a lab document. Essentially, it is correctly positioned that essentially the consent is just another document of a particular format, and having a particular vocabulary and meaning to it in the same way that a medical summary has a particular vocabulary and meaning to it. And if you are interested in finding a medical summary, you know how to find it, and you know how to interpret it when it's there, so if you need a consent document, you know how to find it, and you know how to interpret it. But there are also rules in there that explain that this is a document intended to mediate the access to the other documents.

So when we get done producing the BPPC profile, we come up with these value propositions, what is the usefulness of this particular profile. So within an XDS affinity domain, essentially what it says is the affinity domain should write a set of privacy policies. And this is that first step that we talked about on writing the policy. What does opt in mean to your environment? What does opt out mean to your environment? First is, what are the different policies you're going to have, and what do they mean? Unfortunately, from a BPPC perspective, that is purely an administrative policy-writing task. There is not a connection per se between that language and any other language that BPPC supports.

But what BPPC says is for each one of those that you write, you will assign an OID to, a unique identifier to. In that way, whenever anyone sees this unique identifier, they know that, oh, that is referring to the opt in policy for this health information exchange. So it sets up this concept that because there isn't a technology language back in 2006 or how to encode a policy, we will just simply say that the policies are written and give them a unique identifier. It's an out of band task for how to configure your access control engine to enforce that policy. But clearly you must configure your access control engine to enforce that policy before your access control engine is allowed to say that it understands that policy. So there's a very important part of this, which is the corollary that says, hey, if there is a policy OID, that your access control engine does not understand the meaning for, it's not an enabled policy. So you cannot use the absence of a policy knowledge to enable access by an individual.

This then sets up the option, so you have this set of policies that a health information exchange could publish as these are policies to be used. This would allow the patient to be given the option as to which one of these or which ones of these they would like to authorize, and they would authorize through possibly putting ink onto paper, possibly a digital signature, possibly just a conversation with a clerk, that process part that I talked about at the beginning. Effectively where IHE profile comes in is how do you capture that that act has happened in the very way that how do you capture that the act of a surgery has happened. So you're only capturing the act that the consent has happened. And it explains how you capture that, and we'll get into that.

Here's an example of a policy that could very well be written by a health information exchange, and we actually wrote this one into the profile because we wanted to show you that a doctor wearing a chicken costume, you know, actually if that's your policy, and you have configured somehow to get your access control engines to enforce that policy, and you can possibly get a patient and a doctor to carry out this policy, it's perfectly acceptable from a BPPC perspective. Literally, it is up to the written text to explain what the policy is.

To be more practical, though, we've looked at a bunch of different policies and captured them from many different places. Essentially, these are the kinds of policies that you could support explicit opt in. And you'll notice that you actually have to get rather explicit about this information. You could support an opt

out environment. You could support an opt in environment for certain classes of documents only. You could opt out completely that the documents are not even allowed to be published, saying nothing of whether they're going to be allowed to be used.

You could opt in for emergency access only. Say, hey, look. I don't like being in the health information, but I'm willing to accept that there are going to be emergencies in my life. So these are the kinds of things that somebody could write. And they're simple enough that they could be implemented by just simply knowing that this policy in total has been accepted.

But you'll notice, there's no such thing here as well, yes. I'll opt in, but Dr. Bob, I don't want to see my data. So unfortunately, with basic patient privacy consent, you can't satisfy the needs of some consent requirements for things like saying, hey, Dr. Bob is a neighbor of mine. Or for saying I only want Dr. Joe to see my data. So you start to get into some of these intricacies that become a particular problem.

There certainly is some form of dynamics to it. And if we actually look at the recently published white paper from HHS on consent, they broke this down into there really is, operationally today, five different kinds of consent. I would argue that I think there's a little bit more that they actually describe in that white paper, but we'll go with their five breakdowns.

Surely you can support a no consent environment by just not even using BPPC or using BPPC only for the opt out. Okay. Well, we've got two of their use cases covered. And then clearly the opt in can be supported. There are exceptions, and there are opt in with restrictions. There are some things that you can support. For example, BPPC does support that the consent has a particular effective time, so you can support that this consent is only effective for the next 30 days by specifying the end date. But of course, many of the exceptions are not really going to be easily supported.

Now we're going to get into more about how did we do this? First up, it's really useful to reexamine the characteristics of the CDA document. These are true of all CDA documents. This is some primitive of the structured document standard from HL-7. A CDA document has persistence. Every document is given a unique identifier that identifies that document in perpetuity. They have stewardship that explains who is controlling this document and who authored it and the like. You have this potential for authentication, i.e. where did it come from. Who was the author?

It has a context built into it. Where was this CDA document captured? What is the context of the act that is being captured within this CDA document? So this is important when you indicate that you're capturing it at St. Luke's Hospital for the context of the whole health information exchange. There are some parts of that classic CDA capability that are highly important.

It has the ability to … that it's whole, so you can tell that you have a fragment of a document by noting that it's not a complete XML document. It's got that wholeness concept. Then there is this human readability aspect that CDA has as a tenant, so there is the issue that it's not a binary blob that you absolutely have to have a specialized application to view. In addition, it has specific parts that are intended to be read.

So we really felt that this was an important aspect, so we built off of it. We said, if we just use CDA, we get these free, so to speak. So we then go to the act of consenting, and we say, well, okay. We need to capture in the CDA document that one or more of the affinity domains defined policies have been agreed to. It does support; it's not just one policy, so you can actually set up an environment that might take all of the classifications of sensitive topics and indicate, for each one of those as a policy. Yes, you'll end up

with a couple of dozen policies, but you end up with these fragments of how to deal with these things. You can actually say, I agree to policy one and 12 and 23 and 24.

We indicate in the CDA document when did this consent happen and how long is it good for. As I indicated, for episodic consent, it says this is only good for this surgery episode that I'm going in for. Okay. A surgery, the event, will probably be over in 30 days, so let's do that this consent is only good for 30 days. Past that 30 days, the consent will be not valid. You indicate that it's a template from the BPPC profile, so in the CDA document it declares that it's using the BPPC template.

Potentially, you're using a scanned document, which would be the scan of the ink on paper. Potentially, you're using the digital signature. It could be the patient, but it could also be a guardian. It could also be the clerk, the registration desk clerk, or it could even be the system in the case of Dixie's slide. She showed that the consent could possibly be captured using the classic checkout counter signature-capturing device. Maybe that device itself not only captures the ink signature in wet signature format, but also digitally signs the consent itself.

You then update. You declare the XDS metadata that is a BPPC document. You put into the XDS event code list that the events that are being captured in here is that same list of policy unique identifiers. Then, of course, because it's any document, this consent document itself might be sensitive or fully available. Of course, you also have to, just like any other document, indicate its confidentiality code.

This is the pictorial view of the same thing. The green is the CDA document. It's made up of the CDA header with the patient and the author and the authenticator and the institution and the time of service. The scanned document is an attachment. There the orange is the capturing of the act, which is just simply capturing the policy ID numbers that are being acknowledged. And the red is a digital signature that is signed across the CDA document. Then, of course, the metadata, which indicates that it's a consent document and a digital signature document.

Quickly here, we're reusing CDA release two. We're reusing the XDS scanned documents profile, which is really a profile of PDFA for archive, which is ISO 19005. It's essentially a profile of Adobe PDF that has removed out things that wouldn't likely survive 100-year archives. The XDS document digital signature profile is used for the digital signatures, which is an XML digital signature using the ZATIS profile, and a couple of vocabularies for the purpose of the signature and the time of day of the signature and such. Then, of course, it's got some requirements on XDS, and XDM, and XDR, and XDA on when you see these consent documents, how are you to process them, those kinds of things.

This gets quite into the weeds. I'm not sure if it's too important to get through here much, but essentially you can always indicate within the – this slide is on the document consumer side, so ten years from now, somebody is going in and going to look for a document. Certainly they would be informed by the confidentiality codes on the documents that are available, as to indicate the sensitivity of each of those documents. They would obviously know who they are and who the patient is and what their setting is and how they're going to use it, the purpose of use.

They would apply the access control engine available according to this context information. And they would be able to pull down what are the active consents that are available through a rather simple query that just says please give me the consent documents because they are a type of document, and you can look at the event code list to see which policies have been acknowledged.

There are some special cases that we had to add later with XDM and XDR because originally we were focused only on XDS. In the XDR and XDM environment, they don't really have a declared affinity

domain. So as much as you want to say we know that everybody who's receiving these documents has already agreed to the policies, you have to have some kind of out of band knowledge that who you're sending it to will enforce the policies or that ultimately where the USB memory stick goes is either, A, mediated by the patient themselves who is responsible, or some other mechanism.

Essentially, you should put onto these transactions or media the relevant consents, so as you export a medical summary, what is the relevant consent that should go along with that? Put those onto the media as well. An importer of any documents in these environments has to probably coerce the confidentiality codes, coerce the consent policies into their own understanding. We don't have the algorithm for how to do that, but we explicitly indicate that it's a concern in those environments. It's a risk that needs to be mitigated by policy. Then, of course, things like talking about where the patient is carrying the USB memory risk, there are risks associated with that as well.

This is just kind of a listing of some of the standards that were involved. Then this is my last slide, which is really a look forward, as opposed to what we've been doing. We've been looking at what we wrote, what IHE wrote back four years ago. This is some of the activity that's currently going on and hopefully will produce fruit in the future, and you'll notice that it's quite a few organizations that are working together on this. We have Oasis that have some healthcare specific verticals working on things like the … I don't even have the acronym listed there, but profiles of SAML and XACML and WS Trust for healthcare purposes in cross-enterprise transactions.

We also saw the work in our last meeting from the policy building committee … I don't have that listed here, but I certainly should. HL-7 is working on an update to the CDA with a profile for capturing consent directive that would be able to carry more dynamic attributes, so we could potentially say Dr. Bob is not allowed. They're also working on an ontology for security and privacy. They're working on privacy policy reference catalogs, which is essentially potentially a catalog of privacy policies that would be fully written and be given unique identifiers at the HL-7 level. If any particular exchange wanted to use them, they could just simply use the unique identity value, and because it's a standard policy, people would understand what that policy is.

The fact that it exists doesn't mean that it's endorsed by anybody. The fact that they use it means that they have endorsed it. This concept of creating the catalog of consent policies or consent fragments is, I think, a very neat idea that'll get us somewhere in between.

IHE has got an access control white paper. We're continuing to work on the XUA for SAML and in support of things like including what is the consent that should authorize this transaction. ISO is continuing to work on some. I've got one listed there on purposes of use. At this point, I'll take questions from the committee members, I think.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Okay. This is Dixie. You mentioned earlier that the BPPC ultimately really focused on an acknowledgement of the receipt of privacy practices. If I were to implement BPPC in today's environment, I would probably have a different CDA document for acknowledgement of privacy practices, authorization to participate in research for every kind of purpose. Is that right?

John Moehrke – Interoperability & Security, GE – Principal Engineer
Yes. You would actually have a different policy and a unique identifier for those policies. Those policies don't need to be written in a CDA document. Those would be just written out in whatever human readable form the exchange wants to have. Yes, exactly, the research. I authorize research is certainly a very valid policy. But what does that mean? It authorizes who to do what kind of research?

That's the intricacies that the human readable text needs to convey. Yes, that would be given a unique value for which a CDA document could capture that I, John Moehrke, have agreed to that policy in this particular health information exchange for this period of time. A CDA document only captures the event of me acknowledging that policy.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
I see. Yes. But we know that a hospital, for example, most hospitals have multiple forms that they have patients sign, and translating that over into this kind of an environment, you'd probably have an acknowledgement for each of those forms, I guess.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Yes, exactly.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Are there other questions?

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Yes. John, this is Walter. Hello. Good to hear your voice again. A question that has always been in my mind about this, and some of all the models that have been involved, but this particularly is what are some of the key structural elements that are going to be required to operationalize in a practical way this type of a model? When I say what are the structural elements, I mean it sounds like this requires a central place where the policies are maintained so that organizations can see an OID of a policy and then go and reference it out to some data place that maintains it. What are some of those core structural elements that are needed in order for this to happen?

John Moehrke – Interoperability & Security, GE – Principal Engineer
Well, I think you bring up something that many people do gravitate towards, and first off, IHE wouldn't see that as an interoperability problem because there are plenty of Web servers out there that could easily be used to pull down some text associated with a unique ID value, and it wasn't seen as an interoperability problem.

There are lots of operational things that I kind of alluded to at the beginning, which is, A, you have to write these policies. Yes, you need to publish them. You need to have somebody configure their access control engine to understand what that OID means in their access control engine language. You have to do the typical provisioning that says have I made sure that this system will enforce the policies before you actually authorize them to even communicate at the network level, so way down at the TLS level is where you make those go, no go decisions, and certainly wouldn't authorize a system to connect to health information exchange if they're not willing to enforce the policies that they would have to enforce.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
Who makes that governance decision?

John Moehrke – Interoperability & Security, GE – Principal Engineer
I don't really know if I have that list of primitive that you're talking about.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
But who makes the decision regarding that participation at a human level or at an organization level?

John Moehrke – Interoperability & Security, GE – Principal Engineer

It very much depends on how your health information is organized, your HIO.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
What happens if you go to the next level up, and you just say, we're going to ultimately have this NHIN? How do you adjudicate that, or has that not really been thought through yet?

John Moehrke – Interoperability & Security, GE – Principal Engineer
Well, the NHIN did have to go through that exercise in the absence of some of those things, and I don't want to necessarily speak for that committee, but they took this model. They augmented it slightly because they accepted that BPPC content is one possible way to record an acknowledgement to a consent, but you could also use an XACML blob itself, which would then give them the ability to carry dynamic. But essentially what they had to do was they had to sit in the middle of two organizations, one that might have different policies than the other. Yes, it is a very difficult process, and those on the HIT policy committee, I'm sure, are well aware of just how difficult it is to deal with this human factor of writing policies and promulgating policies, and then doing the post audit that the policies are being enforced. And, if not, what do you do?

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
Well, I don't think that exists yet, does it? This is John Houston. Deven McGraw, do you have thoughts on that?

Deven McGraw - Center for Democracy & Technology – Director
Sorry, John. You should have let me know that you were going to call on me because I stepped aside when you asked your initial question. But I stepped back in right when you said my name.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
I'm sorry, Deven. I'm sorry. I didn't mean to put you on the spot.

Deven McGraw - Center for Democracy & Technology – Director
No, that's okay. But what in particular? Sorry.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
We were just talking about the whole idea of organizationally; the comment was that the HIT Policy Committee was thinking of how ultimately we'd do what amounts to governance, so oversight of who adjudicates these authorizations, I guess, is the best way to put it.

Deven McGraw - Center for Democracy & Technology – Director
Well, I mean, so current law has adjudicatory authority, and that's what we have in the form of consent policy today.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
Right, but the organization that supports that, I guess, is the question I guess I had. Sorry.

Deven McGraw - Center for Democracy & Technology – Director
The organization, like an HIE for example?

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
The Office of Civil Rights enforces it.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security

Who does?

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
The Office of Civil Rights enforces what's contained in HIPAA and ARRA. There are some new policies that are in ARRA that are beyond HIPAA, like for example the patient's right to deny access to a health record if they pay in cash. That's part of the—

Deven McGraw - Center for Democracy & Technology – Director
It's only disclosures to payers. It's not an access denial.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
My apologies. Maybe I'm getting off base. I wasn't really speaking of enforcement per se, but ultimately somebody has to adjudicate some of these requests and how the entities participate in establishing these relationships in which data may get exchanged and consents be passed and things like that. I was just trying to, at a more macro level, try to understand ultimately who has accountability for overseeing all of these organizations.

Deven McGraw - Center for Democracy & Technology – Director
I suspect there'll be multiple layers of that.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
But it really hasn't been defined.

Deven McGraw - Center for Democracy & Technology – Director
The government will play one role, and maybe there might be accrediting bodies on HIEs that will play another. I mean, that's all what's being discussed now.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
Right.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
And even the local hospital; they have local policies.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Yes. I mean, I think one of the realities is that when one looks at policies, one can think of three. I mean, as an organization, one sees three different types of policies. One is policy that I'm required by federal or state law to offer the patient as a consumer preference choice. Number two is some policy that I decide as an organization I'm going to offer above and beyond or outside, not in conflict, but supplement or outside of, above and beyond any federal or state policy or law that I'm going to offer to a consumer to make a choice about privacy preferences. That's those two. Then there's a third one that is the consumer wants to request something that I don't offer by choice or that I'm not required to offer, and so every organization is going to have to deal with those three, and address those three concurrently.

When a CDA document that has a specific policy that are of the category one and category two in my listing, go from one state to another, the recipient in that other state of that CDA will open that and see, oh, in my state, I'm not required to do this, so I'm not necessarily going to – or required to comply with this particularly policy. And I have this other one, so I now have to go to the consumer and offer those, and then tag them into the CDA with OID specific policies. There is that challenge of how all these really work.

That's why I was asking about the implication that there has to be some universal place where these policies get cross-referenced or not necessary a central place. As, John, you point out, there probably are many data places that can handle that type of work. But there has to be some cross-referencing OID system that allows me to know which is this OID and what it means, and in what state or in what jurisdictional context it has to be applied.

John Moehrke – Interoperability & Security, GE – Principal Engineer
I certainly would expect that within any health information exchange, they would operationally put up some way of distributing their policies. And if they choose to use paper, they can choose to use paper. If they choose to use a Web server, they can use a Web server. Essentially that was kind of where I was going with that is that I don't think it's necessarily germane to the interoperability of this particular BPPC profile that the mechanism of communicating the text that's associated with the unique policy ID value is centrally managed. You could put the text into copies of it into every CDA document or every consent document. You have a million patients in your region. You'd have a million copies of it in the system. That's not a very efficient way to do it.

I think the other piece of this is the recognition that NHIN Direct has, which is to recognize that there are point-to-point conversations by which you make all of these decisions completely out of band, and that you have an understanding of the consent environment and understanding of the regulatory environment that you are authorized to communicate from point A specifically to point B, so there are also other simplifications that can be done.

What is the organization? In most cases, there is no organization. It's the sending organization. In an exchange or in a federation of exchanges, there is either an organization or a federation of organizations that need to come to those decisions on what they authorize and how they make sure that those rules are in place.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
That's my point.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
But in the cases of a number of people are exploring the idea of consent registries as well, in which case those would be just as relevant to NHIN Direct as any other exchange because, even with the point-to-point, they could still consult a consent registry.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Sure. And this is something that HITSP in TP30 indicated. The model that I've shown here implies that you are managing for your clinical documents in the same system that you're managing your consent documents. In HITSP TP30 … actually mange your consent documents in an independent consent registry, then your clinical documents. Then you would have a single path for policy communications that is potentially different and managed differently than the path for your consent documents. And that environment that you're talking about, Dixie, would be the consent registry, if you will. I'm not sure how much better that is, but it certainly is a model that isn't explicitly in or out from this profile perspective. The profile didn't want to lay claim to one or the other architecture.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Yes. This is Walter again. It seems to me that we are mixing, in some respects, the uppercase HIE environment with the lower case HIE environment. The uppercase HIE environment meaning a formal, regional, exchange structure in which there is a HIO, and there is a defined set of parameters that organizations that are part of that upper case HIE environment agreed to. But then there is the health

information exchanges in lowercase that happen, which is today the vast majority that happen between two organizations inside a state or across state boundaries in which there's no connection to a regional or even national HIE.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Right.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
And that this model will need to support because I think, for a long while, we're going to see a lot of that happening. Some exchanges happening within a regional, formally established health information exchange in uppercase, and then a lot of exchanges happening point-to-point between organizations outside of an HIE environment.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Yes.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Yes. I think Deven could tell us a lot about those two. She's part of the health information exchange working group of the policy committee as well, which addresses both kinds. Right, Deven?

Deven McGraw - Center for Democracy & Technology – Director
Yes. No, I mean, certainly if you look at where even the discussions are going around the National Health Information Network, the focus is on standing up a set of protocols and policies that will allow for simple, one-to-one exchange that doesn't preclude the use of an HIE uppercase. I like the way you said that, Walter. It sounds so much better than HIE as a noun or HIE as a verb.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
It does.

Deven McGraw - Center for Democracy & Technology – Director
It doesn't preclude the use of an uppercase HIE, but allows exchange to happen in stage one, which is right around the corner of meaningful use, but without needing to have sort of the robustness of an exchange with a capital E in operation in your area. And so this is really sort of a simple set of protocols where the data holder essentially pushes the data either on his or her own initiative or does so in response to a query, but makes the vetting decision about whether that data goes ultimately and, therefore, holds the, in my view, the responsibility of complying with any consent laws that might already be on the books.

Mike DeCarlo – BlueCross BlueShield
This is Mike DeCarlo, Deven. My concern, and I completely agree, but my concern, and I think what we've been playing around or dancing around is whether or not this model supports the continuation of that consent in the recipient, or if the recipient is aware of the limitations of the consent, and then complies with it or re-tasks the data, passes it on again now unknowingly or knowingly, in contradiction to the original consent from the original sender.

Deven McGraw - Center for Democracy & Technology – Director
Well, so I think it depends on where the authority vests for that consent. Let's say if it's within a state and it's a piece of data for which a consent requirement exists, then both discloser and recipient are going to have to abide by it. It's a different set of circumstances if you're crossing state boundaries and the consent requirement applies in one state, but doesn't apply in the other.

Mike DeCarlo – BlueCross BlueShield
I completely agree with that analysis. My question goes to what we've just had presented to us and whether or not the CDA and the consent capability, which John has just described, is robust enough to make that happen.

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
This is Kathleen. I would like to just add here that the model that John has been talking about, and there's been a lot of work around that, and he presented very well that this part of a migration strategy, so the traditional EHRs that may not have the capability to do more robust policy and consent directive support are able to at least deal with the content of the policies and then translate that into their security system.

This same model can be made more mature, as he pointed out in the work that has happened at HL-7 that John has been involved with. In that standard and in the BPPC, it is possible to have technologies that do secure the data for downstream enforcement and requires the recipient to obtain a new license or a new, basically, consent. And this is done transparently to the user, but they're basically going back to a server or the consent directive registry and getting a license to use or permission to use that resource in the way that the policy has described. That's a capability that's used in other industries that might be applied here eventually.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
I appreciate you bringing that up, and it's a good opportunity for me to plug our next session, which is about exactly that. It's about the HL-7 work on the data consent and composite privacy consent directive domain. And it's on, for those of you who are interested, it will be on May 14th, and our speaker is going to be Iona Singare-Rhono. I don't know how to pronounce her last name.

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
Yes. Iona Singare-Rhono.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Okay. She's very good, very good, and so put that on your calendar. Let me see the time.

Iona Singare-Rhono
I have to warn you, I'm on the phone, so not to…

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Okay.

Iona Singare-Rhono
…in my head.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Thank you. I'm glad you dialed in. It's on 2:00 to 4:00. No.

Judy Sparrow – Office of the National Coordinator – Executive Director
No, it's 10:00 to 12:00.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
It's 10:00 to 12:00 on….

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Ten o'clock to 12 o'clock eastern.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
…May 14th.

Judy Sparrow – Office of the National Coordinator – Executive Director
Right.

Deven McGraw - Center for Democracy & Technology – Director
May 4th?

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
May 14th.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
May 14th, 10:00 to 12:00 eastern time.

Deven McGraw - Center for Democracy & Technology – Director
Thank you.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Eastern, yes. It fit in since Kathleen gave the intro, but I don't mean to be truncating the meeting, so back to your comments and questions. Iona, do you have anything to say since you're on?

Iona Singare-Rhono
Actually, I just wanted to clarify perhaps a little aspect of the digital signature that basically the CDA document will not contain a digital signature, but it could indicate that the signature is available, is on file. And, as part of the XDS, for instance, the XDS metadata would include the signatures that are applicable to the document. Also, the XDS metadata may include other kinds of information, and that's where we have to be a little bit careful when exchanging consent directives using BPPC not to expose too much context in the transport metadata in the envelope about the consent itself.

I agree with John. The particular instance of the document is going to acknowledge an opt in or an opt out to a default privacy policy that's applicable in that environment. But you could also say something about who the provider is in the envelope, in the XDS metadata, and you may inadvertently tell someone that this particular consent was signed and the document is available at a substance abuse clinic. In that situation, it is possible to have to leak more information than necessary about the purpose and the context of that consent.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Thank you.

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
I had a slightly different topic I wanted to get back to, if that's okay, and that's the idea, I think John brought this up about what do you do when you're dealing with multiple HIEs, capital HIEs, and sharing records across those various regions or nodes. And also even within an HIE, capital HIE, that it's difficult to know at this juncture whether the type of provider that the BPPC says can see this data if that role means the same for each organization. I think John can speak to this because he's been deeply involved

with the work on trying to standardize the role so that we know the semantics when we talk about a particular provider's role so that we can get at least closer to insuring that each enterprise is interpreting the roles for which access has been denied or permitted in the same way.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Yes. Kathleen, with the BPPC profile, again leveraging basic, it completely begged off on all of that. It basically said the lawyers are going to write the privacy policies in lawyer speak, and it's up to them to get into the legal text the sufficient information, and then it's up to those poor folks who have to configure the access control engines to enforce that policy, to understand the legal text that the lawyers wrote into whatever their access control engine can support. It basically begged off saying we don't have a standard in 2006 or even really much today that bridges those gaps. And that's the topic that we've been working on within Oasis and HL-7, and EXPA has brought together a list of roles specific to your question that are standardized for this kind of a purpose, so you can start to see how I could be using those roles as standardized vocabularies in the context of a dynamic or a consent policy that has blanks in it, if you will.

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
For encoding the value in the code?

John Moehrke – Interoperability & Security, GE – Principal Engineer
Yes.

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
Yes, that's what I meant.

John Moehrke – Interoperability & Security, GE – Principal Engineer
So you fill in the blanks with the vocabulary. But I think, as Iona is going to get to next month, there is absolutely work towards that, but BPPC completely begged off, completely. It just said lawyers will write it, and specialists will configure their access control. And the only thing that will be communicated is that the patient agreed to policy 12. It is basic, but quite honestly, it does deal with those high level opt in, opt out. It just doesn't deal with the opt out with exception or opt in with – you know, so—

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
I think you were talking about how there could be multiple policies, and there's nothing that really precludes you to having a policy that says, in addition to me opting in, I'm excluding providers of this type, and that would be a second policy that the BPPC could be referring to. It's just that it becomes a problem to have multiple scan documents basically guiding what the person is doing who is coding it into the security access control templates, for example.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Yes. Exactly. You could do this. You could actually create 5,000 policies with each possible fragment and say, "I'm authorizing these 60 policies."

Kathleen Connor – Microsoft Health Solutions – Principal Program Manager
I think that just points to the value of the more mature approach that we've all been working on, but the value that the BPPC does bring is for our traditional EHRs that just aren't ready to go there yet. So it's sort of like the accounting of disclosures where I think, in the rule, they said if you had purchased your EHR earlier than a certain date, you were not required to meet the accounting of disclosure until a certain future date. But if you just bought a new EHR, you should be able to do this. Something like that, thanks.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Yes, and I think it was Med Virginia, it might have been, where they actually have a single printed policy that has a bunch of checkboxes, and they implemented that by creating policy fragments for each of the checkboxes. And you then just said here's the sum total of what they checked, so you can do it. But I would shy away if you get more than a dozen policies. It just is going to become more difficult to deal with. The common … explosion is just not worth it, so that's why absolutely I fully endorse what we've been doing in HL-7 and what Iona will be showing next month as the go forward path. The beauty is that we've been working together quite nicely such that if the patient just simply comes in and says, yes, I like the opt in and signs it. It could be a simple BPPC. If they say yes, except then you use the other form of CDA that can support some dynamic attributes.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Great. Are there other comments, questions?

John Moehrke – Interoperability & Security, GE – Principal Engineer
I presume we need to open it up to the public as well.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Is it already open to the public?

Judy Sparrow – Office of the National Coordinator – Executive Director
We need to ask the public if they want to make any comments. Their lines are open, but they can't speak—

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
I see.

Judy Sparrow – Office of the National Coordinator – Executive Director
Yes.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Yes. Let's open it.

Judy Sparrow – Office of the National Coordinator – Executive Director
Okay. If there are no more workgroup calls, operator, could you open the public line, please? Thank you, John, for a great presentation.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Thank you, John. You did a great job. All of you remember that the next session is May the 14th, Friday, May 14th on the HL-7.

Judy Sparrow – Office of the National Coordinator – Executive Director
Right. Any public comments, operator?

Operator
We do not have any questions at this time.

Judy Sparrow – Office of the National Coordinator – Executive Director
Thank you, Dixie. Back to you.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Thank you. Thank you all for dialing in, and thank you, John.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Thank you very much.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Thank you. Thanks, Dixie.

Deven McGraw - Center for Democracy & Technology – Director
Thanks, Dixie.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Thanks, John.

Judy Sparrow – Office of the National Coordinator – Executive Director
Bye-bye.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Bye.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Bye-bye.